CHAPTER

FOUR

HARDWARE

Coordinated computing is concerned with organizing numerous loosely coupled processing agents for coherent action. In this book we emphasize the programming aspects of coordinated computing. However, underlying any computing system is the machinery that performs the computation.

Several multiprocessor (shared-memory) systems have been built (and numerous programming languages designed for them). Siewiorek, Bell, and, Newell's *Computer Structures: Principles and Examples* [Siewiorek 82] contains several articles about such architectures.

Two emerging technologies seem particularly important for the development of coordinated systems. The first is concerned with the inexpensive replication of processing elements—integrated circuit techniques, particularly very large scale integration (VLSI). Perhaps the best survey of this field is Mead and Conway's *Introduction to VLSI Systems* [Mead 80]. VLSI promises to make the price of individual processing elements insignificant, certainly an important development for any system that hopes to exploit many processors.

The study of computer networks seeks optimal ways of connecting computing systems. This field is also called *distributed computing*. Networks connect machines and allow communication. Network technology has matured to the development and product stage. Building a network is no longer a matter of research. Instead, international standards for network communication have been established; several networks are commercially available.

Networks are either local area networks or long haul networks. In a *local area network*, the network manufacturer designs the physical communication channel between computers and can therefore control its performance. Local area networks are used within small areas, such as a single building. *Long haul networks*

Figure 4-1 Network layers.

use common carriers like telephones and satellites for communication. Long haul networks can transfer globe-spanning messages. Of course, the ultimate local area network is many processors on a single chip.

A common network organization is a series of layers or levels, with each layer supported in every network machine. The layers are hierarchical, each one built on the level below it. The bottommost layer of the machines supports the physical signaling and signal recognition. Logically, the n^{th} layer of each machine communicates with the n^{th} layer of the other machines through a series of protocols—rules that describe their conversation. These layers and protocols support not only the immediate task of information transfer, but also serve to route communications, ensure privacy and security, and handle failures.* Important classes of failures include both the failures of individual processing and communication elements and the failures that result from saturating or flooding the network with too many messages. However, the encapsulation provided by multilayered organizations often has its price—a degradation in communication bandwidth through each layer interface. Figure 4-1 illustrates the layering relationship of a computing network.

Ultimately, the goal of a layered network is to hide the details of the processing of each layer from the layers above it. In such an ideal system, communication and processing failures are invisible to the user program. This goal works against the need of some systems for *real-time response*, where the computing system must react to physical phenomena within some time constraint. For example, the disk spinning under the disk head, the ground approaching an airplane, and

* The International Standards Organization (ISO) is in the process of promulgating protocol standards. For example, the X.25 protocol deals with topics such as signaling, transmission errors, addressing, and flow control. Tanenbaum [Tanenbaum 81] is a good source of information about these and other network protocols.

a frustrated user turning off the machine in disgust are all systems that require real-time response (though the time units differ greatly in magnitude). In this book we are concerned both with systems that ignore the problems of response and failure and with languages particularly oriented toward the implementation of real systems.

Good descriptions of the principles and tools of network construction are Ahuja's *Design and Analysis of Computer Communication Networks* [Ahuja 82] and Tanenbaum's *Computer Networks* [Tanenbaum 81].

REFERENCES

[**Ahuja 82**] Ahuja, V., *Design and Analysis of Computer Communication Networks*, McGraw-Hill, New York (1982). Ahuja develops the themes of computer communication networks, starting from the physical connection level, through network design, and on to performance analysis.

[**Mead 80**] Mead, C., and L. Conway, *Introduction to VLSI Systems*, Addison-Wesley, Reading, Massachusetts (1980). This book is a comprehensive discussion of current work on VLSI. Mead and Conway cover material ranging from the physical issues of electron motion through the theoretical issues of circuit complexity.

[**Siewiorek 82**] Siewiorek, D. C., G. Bell, and A. Newell, *Computer Structures: Principles and Examples*, McGraw-Hill, New York (1982). Siewiorek, Bell, and Newell develop the themes of computer architecture both by reprinting important articles and discussing their content.

[**Tanenbaum 81**] Tanenbaum, A. S., *Computer Networks*, Prentice-Hall, Englewood Cliffs, New Jersey (1981). Tanenbaum's book is a comprehensive and well-written discussion of the issues of network design and construction. It is both accessible to the hardware novice and a good reference for the experienced systems architect.